

FAULT TOLERANT COMPUTER, AND DISK MANAGEMENT MECHANISM
AND DISK MANAGEMENT PROGRAM THEREOF

BACKGROUND OF THE INVENTION

5 1. Field of the Invention

 The present invention relates to a lock-step
system fault tolerant computer which processes the same
instruction string in totally the same manner by a
plurality of computing modules in clock synchronization
10 with each other and, more particularly, to a disk
management mechanism of a fault tolerant computer which
facilitates operation required for multiplexing
setting/restoration of a disk.

2. Description of the Related Art

15 In many of conventional fault tolerant computers
of this kind, the disk multiplexing function is realized
by software for the purpose of cost cutting.

 A fault tolerant computer which realizes disk
duplexing by two storage devices for storing an
20 operating system, a user program and user data, for
example, is provided with an access path duplexing
function of making two or more access paths provided for
each of the two storage devices be seen as one from the
operating system and a disk duplexing function of making
25 the two storage devices be recognized as one virtual
storage device by the operating system, which functions
are realized by software for the purpose of cost

reduction.

When a fault occurs such as a failure of a storage device, a virtual storage device will be considered as a single point for the fault, so that because of characteristics of a fault tolerant computer, it is necessary to quickly cut off the storage device developing the fault and integrate a normal storage device to again conduct duplexing of a disk.

In a case where an end user conducts disk multiplexing setting or restoration operation by himself/herself in a fault tolerant computer which realizes a disk duplexing function by software, the user needs to execute complicated operation requiring a broader technical knowledge at the time of cutting off a storage device having a failure and integration of a device.

As described above, when an end user conducts cut-off of a storage device developing a failure and integration of a device in a conventional fault tolerant computer which realizes a disk duplexing function by software, as compared with a case where the function is realized by hardware, more complicated operation requiring a broader technical knowledge should be conducted to make it extremely difficult for the end user to conduct the operation in question by himself/herself.

Therefore, because of difficulty of replacement

of a disk (storage device) developing a fault by an end user by himself/herself, a large MTBF (Mean Time Between Failure: a mean time from a failure occurring in a computer system until when a next failure occurs) which is a characteristic of a fault tolerant computer is reduced to result in preventing the fault tolerant computer to accomplish its own original object.

In other words, a fault tolerant computer realizing a function for duplexing a disk by software for the purpose of cost reduction has a problem that operability in disk multiplexing setting and restoration is degraded to lose the feature of the fault tolerant computer.

SUMMARY OF THE INVENTION

An object of the present invention is to provide a disk management mechanism enabling an end user to conduct operation for disk multiplexing setting/restoration with simple operation without requiring a special technical knowledge when a fault such as a failure of a storage device occurs in a fault tolerant computer.

According to the first aspect of the invention, a fault tolerant computer having a disk multiplexing mechanism which multiplexes a plurality of storage devices and an access path multiplexing mechanism which sets and multiplexes a plurality of access paths for the

plurality of storage devices, comprising a disk management mechanism which inputs, when a fault such as a failure of the storage device occurs, physical position information of the storage device and operation contents related to the storage device in question to instruct the disk multiplexing mechanism on restoration operation including cut-off and integration operation of the storage device.

According to another aspect of the invention, a disk management mechanism of a fault tolerant computer having a disk multiplexing mechanism which multiplexes a plurality of storage devices and an access path multiplexing mechanism which sets and multiplexes a plurality of access paths for the plurality of storage devices, wherein when a fault such as a failure of the storage device occurs, physical position information of the storage device and operation contents related to the storage device in question are input to instruct the disk multiplexing mechanism on restoration operation including cut-off and integration operation of the storage device.

According to another aspect of the invention, a disk management program of a fault tolerant computer having a disk multiplexing mechanism which multiplexes a plurality of storage devices and an access path multiplexing mechanism which sets and multiplexes a plurality of access paths for the plurality of storage

devices, which executes, when a fault such as a failure of the storage device occurs, a function of instructing the disk multiplexing mechanism on restoration operation including cut-off and integration operation of the storage device by inputting physical position information of the storage device and operation contents related to the storage device in question.

Other objects, features and advantages of the present invention will become clear from the detailed description given herebelow.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood more fully from the detailed description given herebelow and from the accompanying drawings of the preferred embodiment of the invention, which, however, should not be taken to be limitative to the invention, but are for explanation and understanding only.

In the drawings:

Fig. 1 is a block diagram showing an entire structure of a fault tolerant computer according to an embodiment of the present invention;

Fig. 2 is a block diagram showing a structure of a disk management mechanism of the fault tolerant computer according to the embodiment of the present invention;

Fig. 3 is a diagram for use in explaining the

contents of a physical position access path conversion
DB of the disk management mechanism shown in Fig. 2; and

Fig. 4 is a sequence diagram for use in
explaining operation of the disk management mechanism in
the fault tolerant computer according to the embodiment
of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

The preferred embodiment of the present invention
will be discussed hereinafter in detail with reference
to the accompanying drawings. In the following
description, numerous specific details are set forth in
order to provide a thorough understanding of the present
invention. It will be obvious, however, to those skilled
in the art that the present invention may be practiced
without these specific details. In other instance, well-
known structures are not shown in detail in order to
unnecessary obscure the present invention.

Embodiment of the present invention will be
described in detail with reference to the drawings in
the following.

Fig. 1 shows an entire structure of a fault
tolerant computer according to an embodiment to which
the present invention is applied.

With reference to Fig. 1, a fault tolerant
computer 10 according to the present embodiment includes
a plurality of computing modules 11 and 12, each of

which computing modules 11 and 12 processes the same instruction string in clock synchronization with each other and compares a processing result of each computing module to enable, even when one computing module
5 develops a fault, the processing to be continued by the remaining computing module.

The computing modules 11 and 12 include a plurality of processors 101 and 102, 201 and 202, processor external buses 401 and 402 and memories 301
10 and 302, respectively.

The fault tolerant server 10 further includes two storage devices 21 and 22 for storing an operating system, a user program or user data, access path duplexing mechanisms 31 and 32 for bundling a plurality
15 of access paths to the two storage devices 21 and 22 into one, a disk duplexing mechanism 40 for making the storage devices 21 and 22 be seen as one from the operating system or the user program through the access path duplexing mechanisms 31 and 32, and a disk
20 management mechanism 50 for accessing the disk duplexing mechanism 40 to provide a simple interface to an end user in disk duplexing setting/restoration operation conducted at the time of restoration or addition of a new storage device when duplexing of a disk is hindered
25 due to a failure of a storage device or a failure in an access path. In Fig. 1, illustration is made only of a characteristic part of the structure of the present

embodiment and that of the remaining common part is omitted.

The storage devices 21 and 22 store an operating system, a user program and user data. As a feature of the fault tolerant computer 10, two or more access paths to the storage devices 21 and 22 are provided and for making these access paths be seen as one access path from the operating system, the access path duplexing mechanisms 31 and 32 are provided.

Moreover, although the storage devices 21 and 22 are seen as a total of two storage devices through the access path duplexing mechanisms 31 and 32, the disk duplexing mechanism 40 for duplexing the two storage devices 21 and 22 makes the storage devices 21 and 22 be recognized as one virtual storage device by the operating system.

On the other hand, when a fault such as a failure of the storage devices 21 and 22 occurs, the virtual storage device is considered to be a single point for the fault, so that the storage device developing the fault should be quickly replaced with a normal storage device to again conduct duplexing of the disk because of the characteristics of the fault tolerant computer.

Here, many of low-cost fault tolerant computers implement the disk duplexing mechanism 40 by software and many of the disk duplexing mechanisms 40 realized by software accordingly need complicated processing

requiring a broader technical knowledge for cutting off
a storage device developing a failure and integration of
a new device, whereby an end user will have an extreme
difficulty in conducting the relevant processing by
5 himself/herself.

Under these circumstances, the present embodiment
is designed such that the disk management mechanism 50
has an interface with the disk duplexing mechanism 40 to
take out actual access path information from the access
10 path duplexing mechanisms 31 and 32, thereby mapping
information about an access path to a storage device
developing a failure and access path information
obtained from the access path duplexing mechanism,
specify a storage device managed by the disk duplexing
15 mechanism 40 and instruct the disk duplexing mechanism
40 to cut off the storage device in question or
integrate a new device.

This arrangement enables an end user to execute
replacement of the storage devices 21 and 22 with ease
20 only by grasping a physical position of the storage
devices 21 and 22.

With reference to Fig. 2, the disk management
mechanism 50 includes an access path duplexing mechanism
access unit 51, a disk duplexing mechanism access unit
25 52, an interface supply unit 53 and a physical position
access path conversion DB (data base) 54.

The access path duplexing mechanism access unit

51 accesses the access path duplexing mechanisms 31 and 32 to obtain information about mapping between the information about the access paths to the storage devices 21 and 22 and access path information duplexed
5 by the access path duplexing mechanisms 31 and 32 which is to be operated by the disk duplexing mechanism 40.

The disk duplexing mechanism access unit 52 accesses and instructs the disk duplexing mechanism 40 on the access path information and a kind of operation
10 (cut-off or integration) to realize cut-off or integration of a specific storage device from or into a virtual storage device.

The interface supply unit 53 obtains access path information of a storage device from the physical
15 position access path conversion DB 54 based on physical position information of the storage device applied by an end user, obtains the access path information and a kind of operation for the disk duplexing mechanism 40 applied by the end user and uses the access path duplexing
20 mechanism access unit 51 and the disk duplexing mechanism access unit 52 to provide the end user with a simple interface.

The physical position access path conversion DB 54, as shown in Fig. 3, stores physical position
25 information indicative of the storage devices 21 and 22 and access path information for the storage devices 21 and 22 so as to correspond with each other.

Next, operation of the present embodiment will be detailed with reference to Fig. 2 and the sequence diagram shown in Fig. 4. Assume here, as illustrated in Fig. 2, that access paths which are served for the disk duplexing mechanism 40 to discriminate and control the storage devices 21 and 22 are access paths A and B and that access paths provided by the access path duplexing mechanisms 31 and 32 for the storage devices 21 and 22 are access paths A1, A2 and access paths B1 and B2.

First, an end user applies physical position information of a storage device to be operated (to designate the storage device 21 or the storage device 22) and operation contents (to designate cut-off or integration) to the interface supply unit 53.

Next, the interface supply unit 53 having received the above-described information accesses the physical position access path conversion DB 54 to obtain access path information of the storage device in question from the physical position information (Sequence A in Fig. 4). In a case, for example, where the storage device 21 develops a failure and the storage device 21 is designated as physical position information in order to conduct cut-off of the device or integration, obtained from the physical position access path conversion DB 54 shown in Fig. 3 is information of (access path A - access path A1) and (access path A - access path A2) as access path information corresponding

to the storage device 21.

The interface supply unit 53 having obtained the above-described access path information transmits the access path information to the access path duplexing mechanisms 31 and 32 through the access path duplexing mechanism access unit 51.

The access path duplexing mechanisms 31 and 32 having obtained the access path information refer to self-managed access path information and when the transmitted access path information exists, reply to the interface supply unit 53 through the access path duplexing mechanism access unit 51 with access path information composed of a virtual access path which is a path obtained by considering two access paths duplexed by the access path duplexing mechanism 31 or 32 as one access path (Sequence B in Fig. 4).

Here, a virtual access path is a path served for the disk duplexing mechanism 40 to discriminate the storage devices 21 and 22 without using the access paths A1, A2, B1 and B2 and in a case of the access path duplexing mechanism 31, it makes a reply with the access paths A1 and A2 as one virtual access path A which is the same as the access path A.

The disk duplexing mechanism 40 only controls duplexing for the storage devices 21 and 22 through the access paths A and B and grasps nothing about the access paths A1, A2, B1 and B2 provided by the access path

duplexing mechanisms 31 and 32. Therefore, for the disk duplexing mechanism 40 to control the storage devices 21 and 22 without using the access paths A1, A2, B1 and B2, such a virtual access path as described above is used.

5 Furthermore, the interface supply unit 53 transmits access path information for the obtained virtual access path (access path A in a case of the storage device 21) and the operation contents applied by the end user to the disk duplexing mechanism 40 through the disk duplexing mechanism access unit 52. With
10 respect to the designated access path information, the disk duplexing mechanism 40 executes operation designated by the operation contents and replies to the interface supply unit 53 with the operation results
15 through the disk duplexing mechanism access unit 52 (Sequence C in Fig. 4). As a result, the interface supply unit 53 notifies the end user of the operation result.

20 When a fault such as a failure of the storage devices 21 and 22 occurs, the foregoing operation enables the end user to instruct the disk duplexing mechanism 40 to cut off or integrate the storage device only by simple operation of inputting physical position information which designates a storage device and
25 operation contents related to the storage device, thereby allowing operation required for duplexing setting/restoration to be conducted without a special

technical knowledge.

In the fault tolerant computer of the present invention, the function of each unit which executes the disk management function can be realized not only by
5 hardware but also by software by the execution, on a CPU, of a disk management program 100 which executes the function of each of the above-described units. The disk management program 100 is stored in a recording medium such as a magnetic disk or a semiconductor memory and
10 loaded into a memory of the CPU from the recording medium and executed by the CPU to realize each of the above-described functions.

Although the present invention has been described with respect to the preferred embodiment in the
15 foregoing, the present invention is not always limited to the above-described embodiment and can be realized in various forms within the scope of its technical idea.

While the embodiment has been described with respect to a case where two storage devices are duplexed
20 by each disk duplexing mechanism 40, it is apparent that the present invention is similarly applicable to a case where three or more storage devices are multiplexed by a disk multiplexing mechanism. Also as to an access path duplexing mechanism, application of the present
25 invention is not limited to duplexing but is possible to a case where three or more access paths are provided by an access path multiplexing mechanism.

As described in the foregoing, when a fault such as a failure of a storage device occurs, the present invention enables the disk multiplexing mechanism to be instructed on cut-off or integration of a storage device only by simple operation of inputting physical position information which designates a storage device and operation contents of the storage device, whereby an end user is allowed to conduct operation required for multiplexing setting/restoration by extremely simple operation without grasping internal access path information and without having a special technical knowledge.

Although the invention has been illustrated and described with respect to exemplary embodiment thereof, it should be understood by those skilled in the art that the foregoing and various other changes, omissions and additions may be made therein and thereto, without departing from the spirit and scope of the present invention. Therefore, the present invention should not be understood as limited to the specific embodiment set out above but to include all possible embodiments which can be embodied within a scope encompassed and equivalents thereof with respect to the feature set out in the appended claims.